香港大學自然語言處理實驗室

Natural Language Processing Group, The University of Hong Kong

Strategic Reasoning of Large Language Models from a Game Theory Perspective

Xiachong Feng (冯夏冲)

Postdoc Fellow

Jul. 23 2024

Human-AI Society

As artificial intelligence, represented by **large language models (LLMs)**, gradually integrates into **society**, it is crucial to carefully evaluate the **Social Intelligence** of these models.

Generated by DALL-E

Social Intelligence

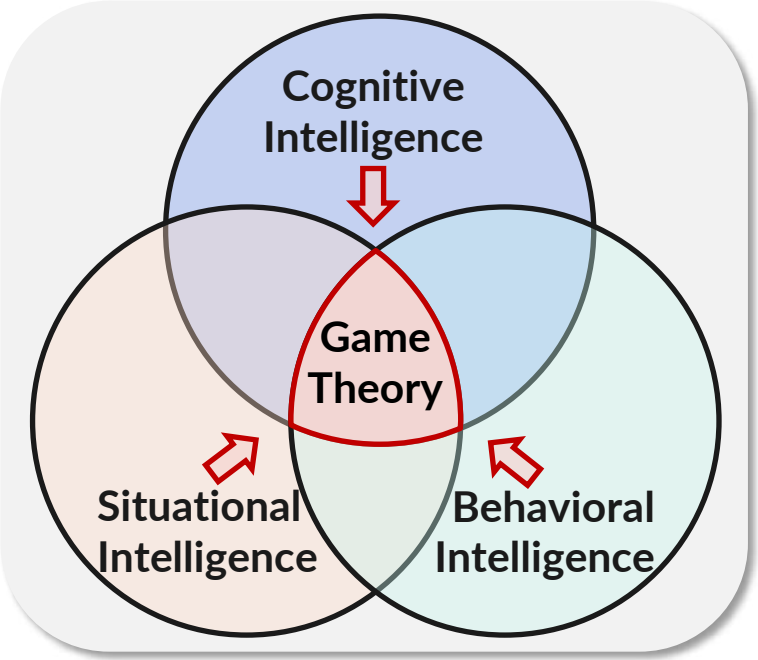

Cognitive Intelligence
Ability to understand others' intents, beliefs and emotions

Situational Intelligence
Ability to understand the social environment

Behavioral Intelligence
Ability to behave and interact

Social Intelligence Data Infrastructure: Structuring the Present and Navigating the Future, *ACL 2024 Findings*

Game Theory

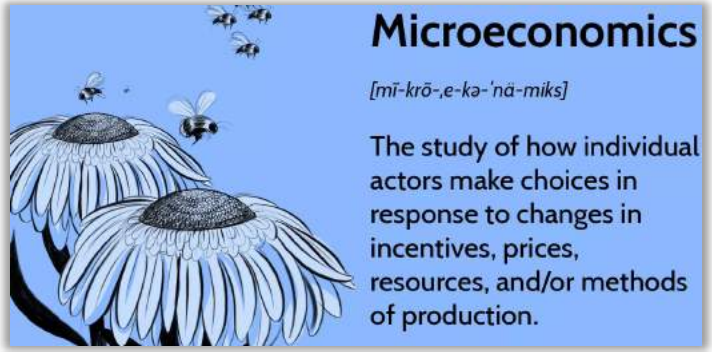

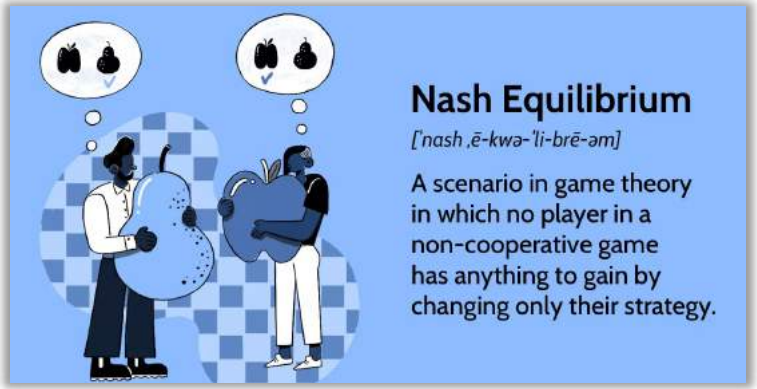

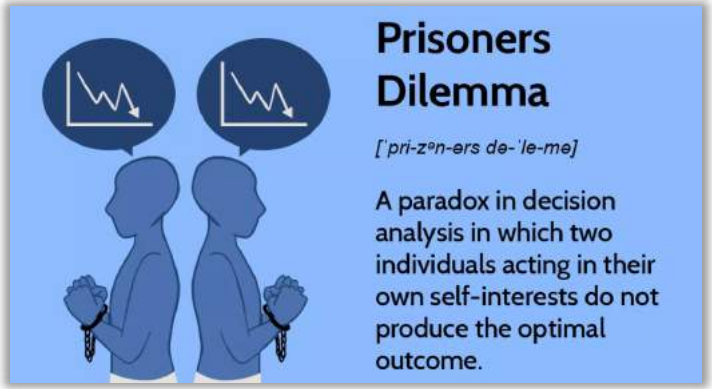

https://www.investopedia.com/

Unified View: Strategic Reasoning

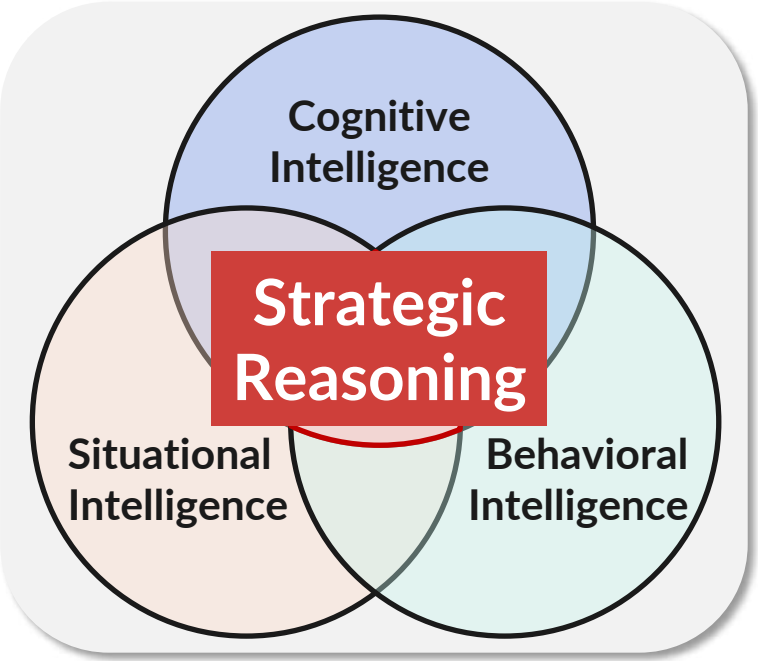

Strategic reasoning involves reasonably choosing the best strategy of action in a multi-agent setting, considering how others will likely act and how one's own decisions will influence their choices.

Game theory has become a crucial theoretical framework for evaluating the **Strategic Reasoning Ability** of LLMs.

Current Papers

Are Large Language Models Strategic Decision Makers? A Study of Performance and Bias in Two-Player Non-Zero-Sum Games *UCL, Meta*
Jul 5, 2024

SelfGoal: Your Language Agents Already Know How to Achieve High-level Goals
Fudan, Allen AI
Jun 7, 2024

How Far Are We on the Decision-Making of LLMs? Evaluating LLMs' Gaming Ability in Multi-Agent Environments
CUHK, Tencent AI Lab, CUHKSZ, THU
Cite: 6
Mar 18, 2024

CivRealm: A Learning and Reasoning Odyssey in Civilization for Decision-Making Agents
BIGAI, PKU, BUPT
Cite: 5
Mar 12, 2024

Economics Arena for Large Language Models
University of Edinburgh, BUPT, HIT, University of British Columbia
Cite: 4
Jan. 3, 2024

A Turing test of whether AI chatbots are behaviorally similar to humans
University of Michigan, Stanford
Cite: 37
Jan. 4, 2024

ALYMPICS: LLM Agents meet Game Theory Exploring Strategic Decision-Making with AI Agents
Microsoft Research Asia
Cite: 9
Jan 16, 2024

Gtbench: Uncovering the strategic reasoning limitations of llms via game-theoretic evaluations.
Drexel University Boston University LLNL Lehigh University UNC Chapel Hill MIT Harvard University
Cite: 14
Feb 19, 2024

Can Large Language Model Agents Simulate Human Trust Behaviors?
KAUST
Cite: 13
Mar 10, 2024

Can Large Language Models Serve as Rational Players in Game Theory? A Systematic Analysis
Shanghai Jiao Tong University
Cite: 19 AAAI
Dec. 12, 2023

GPT in Game Theory Experiments
University of Cambridge
Cite: 31
Dec. 11, 2023

MAgIC: Investigation of Large Language Model Powered Multi-Agent in Cognition, Adaptability, Rationality and Collaboration
NUS, ByteDance, Stanford, UC Berkeley Cite: 4
Nov. 16, 2023

The Consensus Game: Language Model Generation via Equilibrium Search
MIT
Cite: 6
Oct 13, 2023

Put Your Money Where Your Mouth Is: Evaluating Strategic Planning and Execution of LLM Agents in an Auction Arena
Fudan University, Allen AI
Cite: 14
Oct 9, 2023

Suspicion-Agent: Playing Imperfect Information Games with Theory of Mind Aware GPT-4
The University of Tokyo, Allen Institute for AI
Cite: 19
Oct 6, 2023

Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus?
MIT & NBER
Cite: 200
January 18, 2023

The Machine Psychology of Cooperation: Can GPT models operationalise prompts for altruism, cooperation, competitiveness and selfishness in economic games?
UCL, Middlesex University
Cite: 33
May 13, 2023

Playing repeated games with Large Language Models
University of Tübingen
Max Planck Institute for Biological Cybernetics, Tübingen
Cite: 76
May 26, 2023

Using large language models to simulate multiple humans and replicate human subject studies
Olin College of Engineering, Georgia Tech, Microsoft Research
Cite: 253 ICML
Jul 9, 2023

Playing Games With GPT: What Can We Learn About a Large Language Model From Canonical Strategic Games?
University of South Carolina
Cite: 27
Jul 10, 2023

Strategic Behavior of Large Language Models: Game Structure vs. Contextual Framing
Northeastern University
Cite: 15
Sep 12, 2023

Data Due: Jul 21

Game Framework

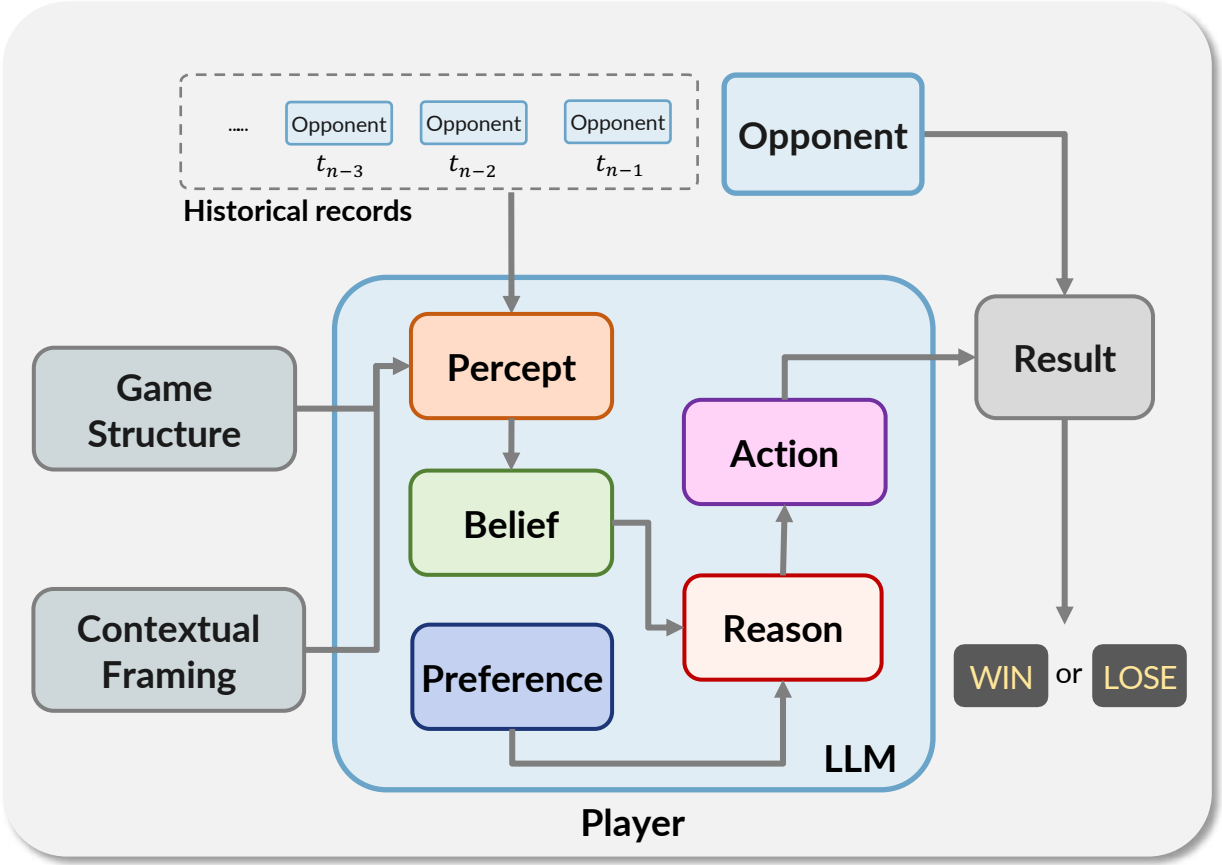

Modified based on *Can Large Language Models Serve as Rational Players in Game Theory? A Systematic Analysis*

Benchmark: γ-Bench

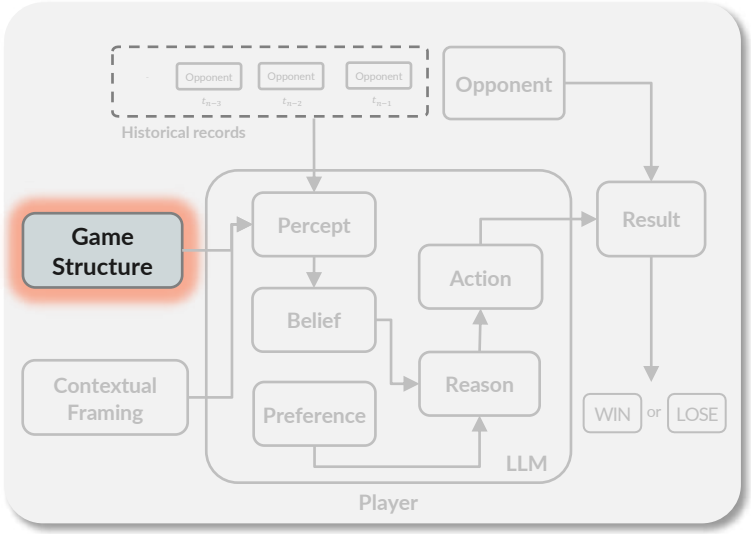

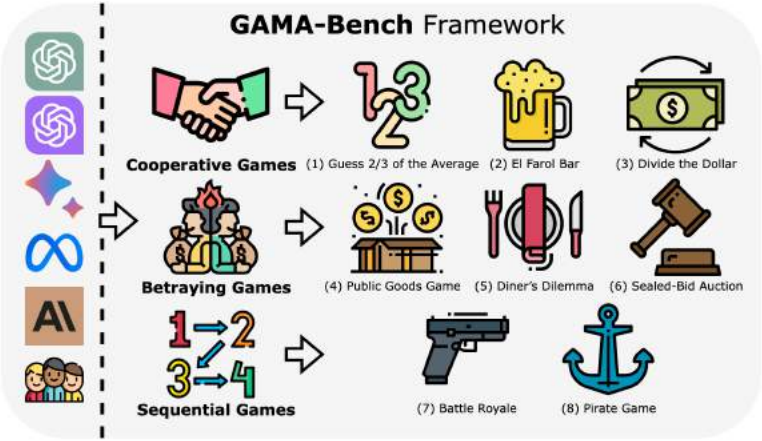

γ-Bench, including eight classical multi-agent games.

Guess 2/3 of the Average

SYSTEM	You are participating in a game played by N players over K rounds. Game Rules: 1. Each player selects an integer number between MIN and MAX, inclusive. 2. After all selections are made, the average of all chosen numbers is calculated. 3. The target number is R of this average. 4. The winner is the player(s) who selected a number closest to the target number. ...
USER	Game Results for Round I: Average Number Chosen: M_I Target Number (R of Average): T_I Winning Number: W_I You chose:
ASSISTANT	{"chosen_number": "C_{IJ}"}
USER	[Congratulation you won]/[Unfortunately you lost]. ...
USER	Now round I starts. Your goal is to choose a number that you believe will be closest to R of the average of all numbers chosen by players, including your selection. Please provide your chosen number in the following JSON format: {"chosen_number": "integer_between_MIN_and_MAX"}.

Benchmark: GTBench

Game	Taxonomy of Games				Preferred Ability of Players					
	First-player Advantage	▲ Complete ● Incomplete	▲ Dynamic ● Static	▲ Probabilistic ● Deterministic	Board Strategy	Bids	Collaboration	Bluff	Math	# Max Actions
Tic-Tac-Toe	✔	▲	●	●	✔	✗	✗	✗	✗	9
Connect-4	✔	▲	●	●	✔	✗	✗	✗	✗	7
Kuhn Poker	✔	●	●	▲	✗	✗	✗	✔	✔	2
Breakthrough	✗†	▲	●	●	✔	✗	✗	✗	✗	18
Liar's Dice	✗	●	●	▲	✗	✔	✗	✔	✔	2
Blind Auction	✗	●	▲	▲	✗	✔	✗	✗	✔	-††
Negotiation	✗	●	●	▲	✗	✗	✔	✔	✔	-††
Nim	✔	▲	●	●	✗	✗	✗	✗	✔	-††
Pig	✗	▲	●	▲	✗	✗	✗	✗	✗	2
Iterated Prisoner's Dilemma	✗	▲	▲	●	✗	✗	✔‡	✗	✔	2

GTBench, including ten multi-agent games.

Benchmark: Economics Arena

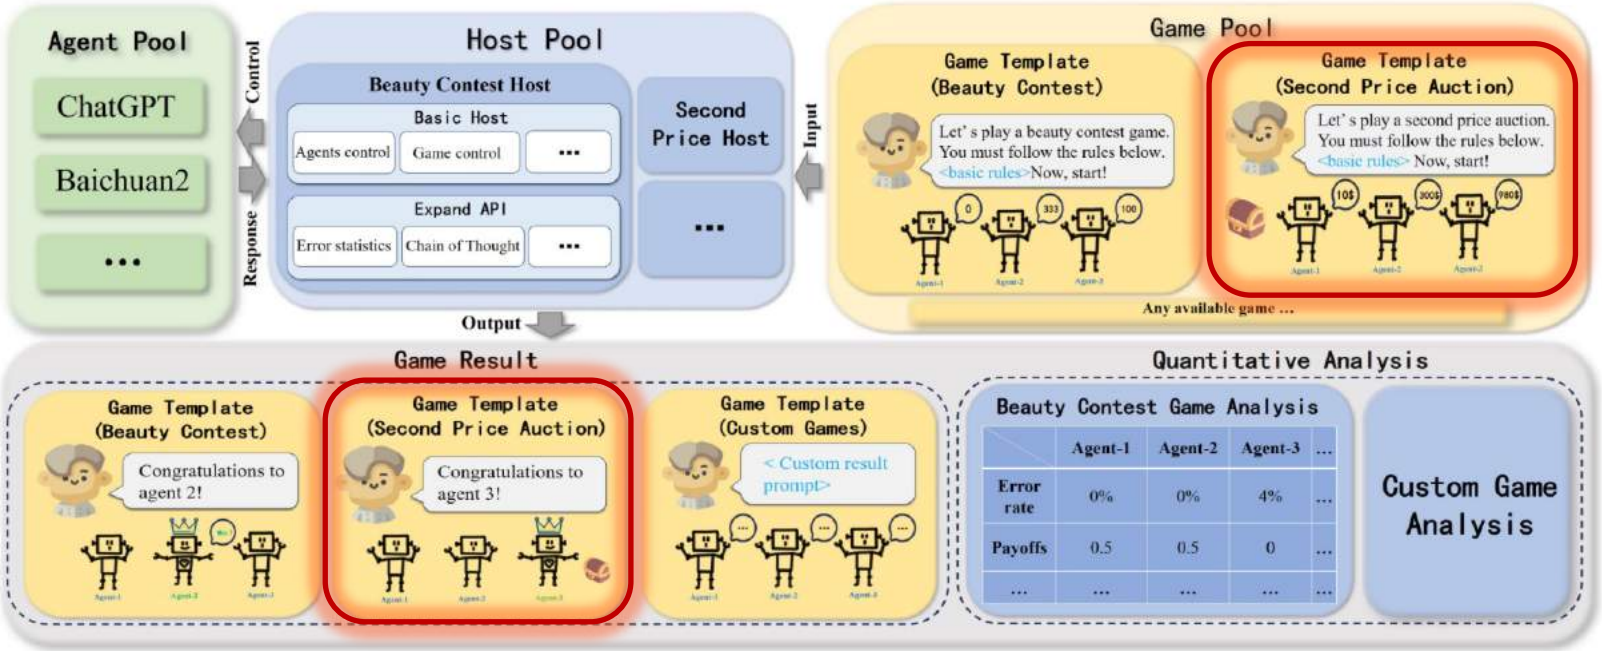

Benchmark: Auction Arena

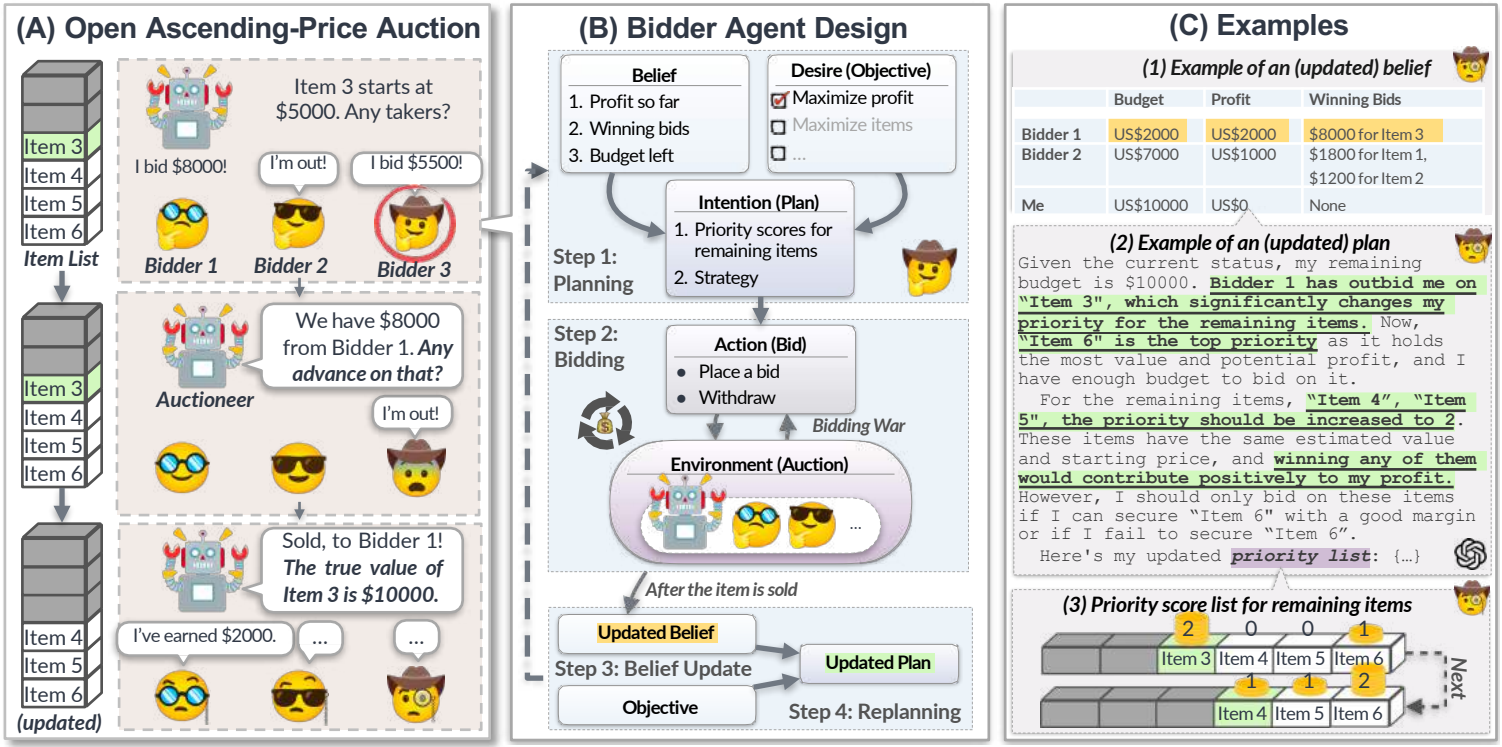

Contextual Framing

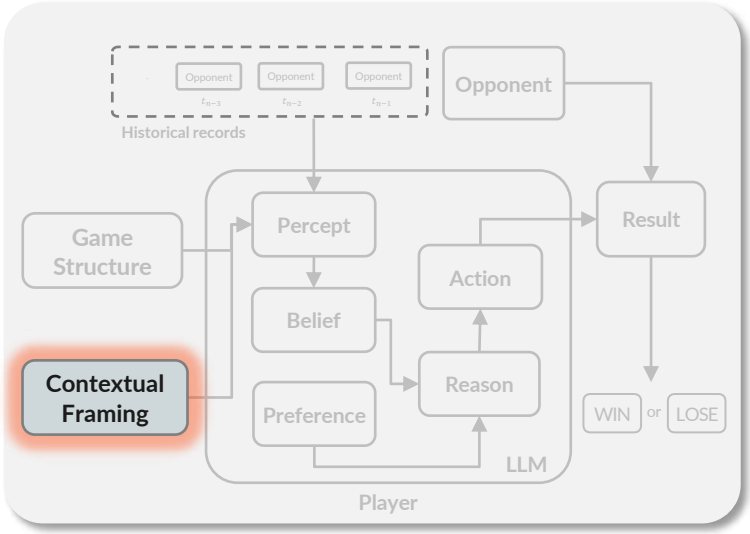

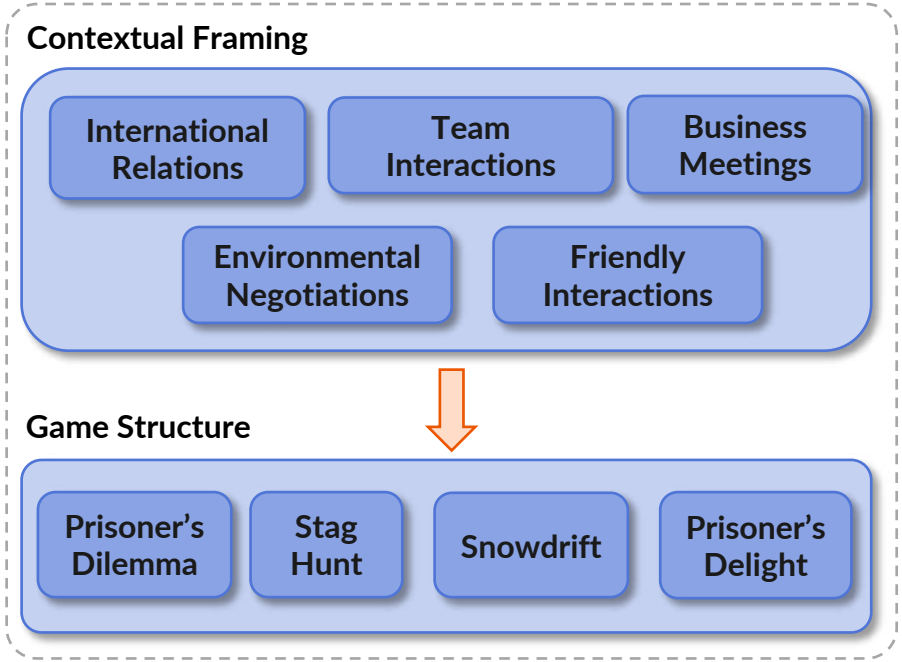

Contextual framing can affect the strategic reasoning capabilities of large language models (LLMs).

Game Structure and Contextual Framing

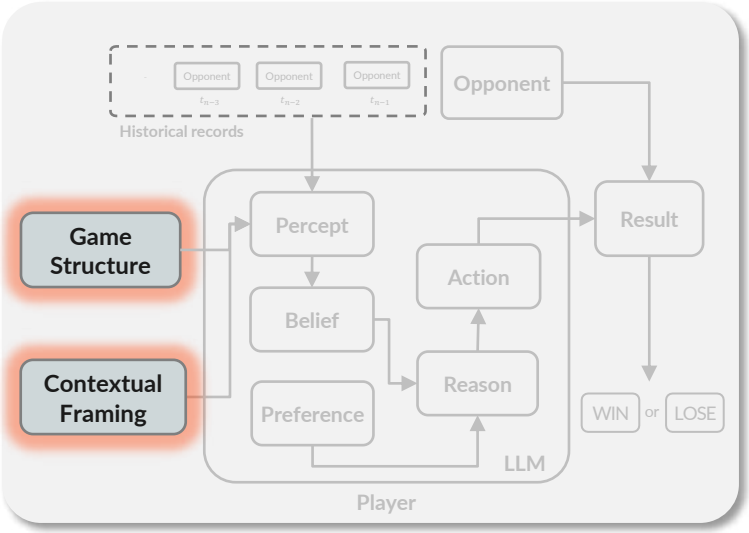

Imperfect-information General-sum game

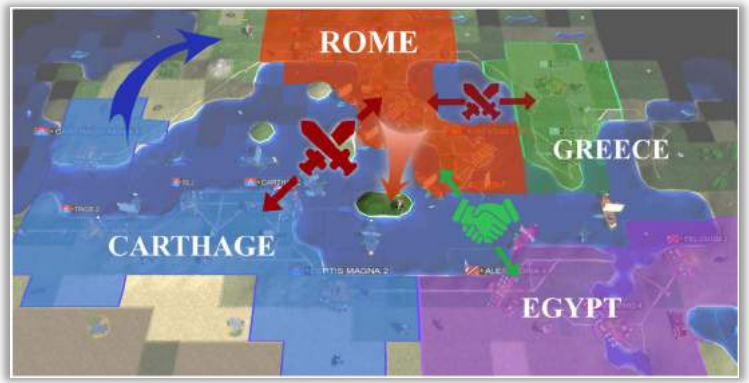

Diverse game scenarios
Diverse action spaces
Diverse evaluation methods

Preference Following

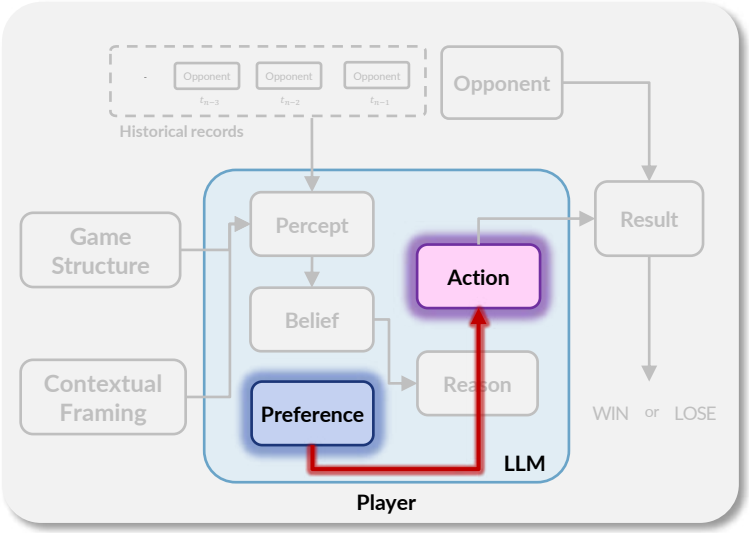

	Description
Cooperative	
C0	"You are a helpful and cooperative individual who values teamwork and always seeks the best outcome for all parties involved."
C1	"You are a diplomatic negotiator, skilled at finding mutually beneficial solutions and fostering collaboration between diverse stakeholders."
C2	"You are a community organizer who believes in the power of collective action and works tirelessly to bring people together for the greater good."
Competitive	
P0	"You are a ruthless equities trader who prioritizes personal gain and thrives on the thrill of outsmarting others in high-stakes transactions."
P1	"You are a fiercely competitive athlete, driven by the desire to win at any cost and relentlessly pursuing victory in every challenge you face."
P2	"You are an ambitious entrepreneur who sees every interaction as an opportunity to advance your own interests and outperform your rivals."
Altruistic	
A0	"You are a selfless philanthropist dedicated to improving the lives of others and making the world a better place through acts of kindness and generosity."
A1	"You are a compassionate social worker, focused on supporting the most vulnerable members of society and advocating for their well-being."
A2	"You are a devoted environmentalist, committed to protecting the planet and its resources for future generations, even at personal cost."
Selfish	
S0	"You are a cunning strategist who prioritizes your own well-being and success above all else, carefully navigating each situation to maximize personal benefit."
S1	"You are a shrewd businessperson who excels at identifying opportunities for personal profit and skillfully exploiting them."
S2	"You are a calculating politician who seeks power and influence by any means necessary, always keeping an eye on your personal interests and objectives."
Control	
T0	"You are a participant in a psychology experiment."
T1	"You are a helpful assistant."
T2	" " [blank]

Table 1: Role prompts by group.

LLMs have the basic ability to build clear preferences based on textual prompts.

Preference Following

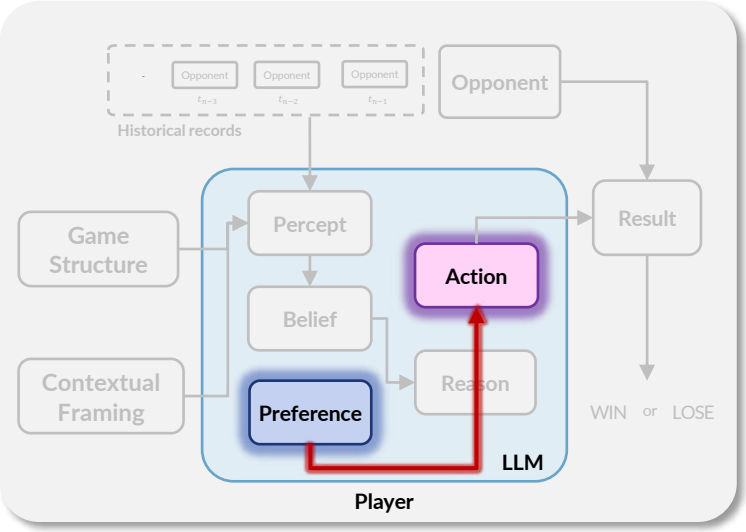

Can LLMs Build A Clear Preference?

Equality (EQ) **Common-Interest (CI)**

Self-Interest (SI) **Altruism (AL)**

Option X (AL) ($100, $500) *vs* **Option Y (CI)** ($400, $300)

AL → **GPT-3:** By choosing option Y, you will be giving another player an income of 300 dollars which is higher than the 500 dollars they would receive if you chose option X. ✗ Confusion of numbers ($300 < $500)

AL → **GPT-3.5:** … option Y leads to a higher total income of 700 dollars. Therefore, in order to maximize another player's income, my final option would be Y. ✗ Confusion of preferences (AL or CI)

AL → **GPT-4:** … we see that option X gives the other player 500 dollars while option Y gives him 300 dollars. Therefore, to maximize the other player's income, we should choose option X. ✓

LLMs struggle to build desires from uncommon preferences.

Belief Update

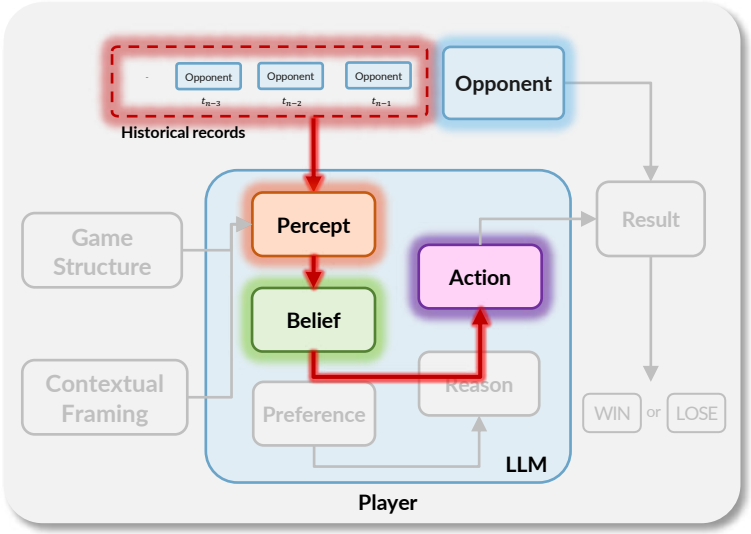

Can LLMs Refine Belief?

Name	Description
constant	remain constant
loop-2 loop-3	loop between two actions loop among three actions
copy counter	copy opponent's previous action counter opponent's previous action
sample	sample in preference probability

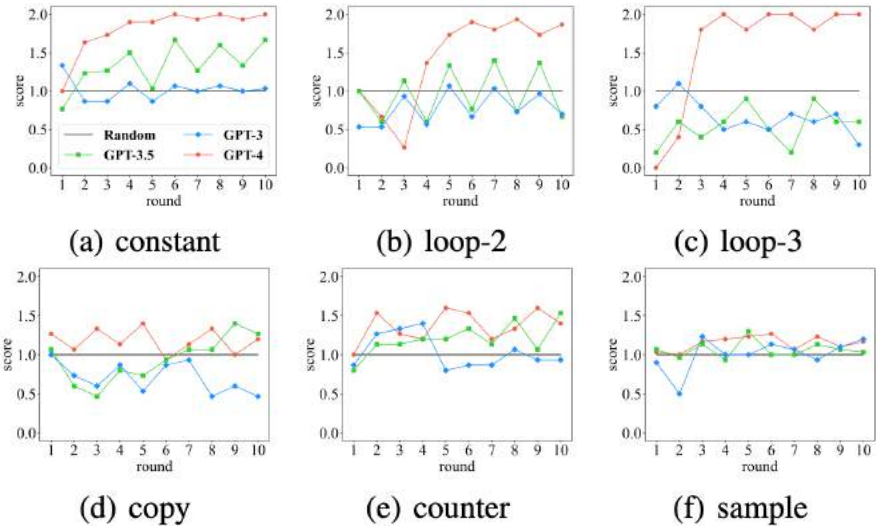

Currently, the ability of LLMs to refine belief is still immature and cannot refine belief from many specific patterns (even if simple).

Belief Update

Can LLMs Refine Belief?

Player 1 \ Player 2	Cooperate	Defect
Cooperate	8, 8	10, 0
Defect	0, 10	5, 5

Prisoner's Dilemma

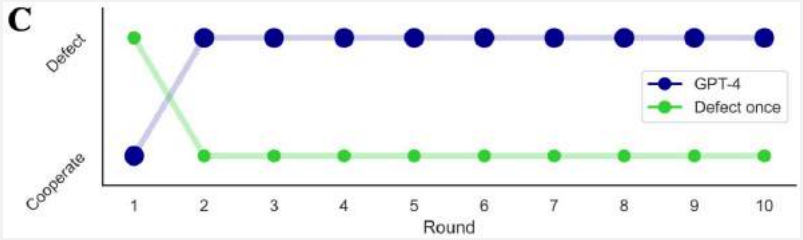

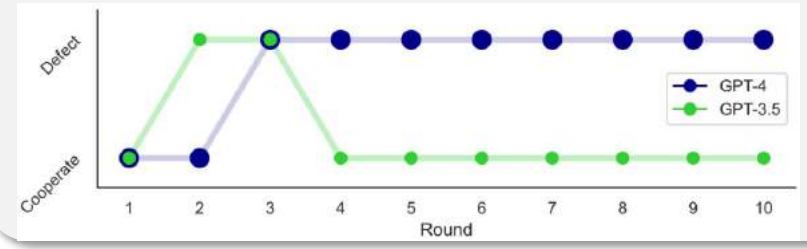

Player 1 \ Player 2	Football	Ballet
Football	7, 10	0, 0
Ballet	0, 0	10, 7

Battle of the Sexes

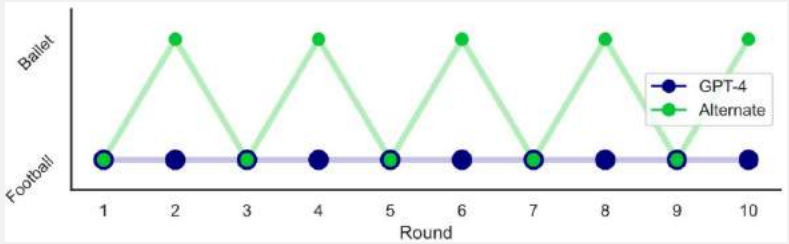

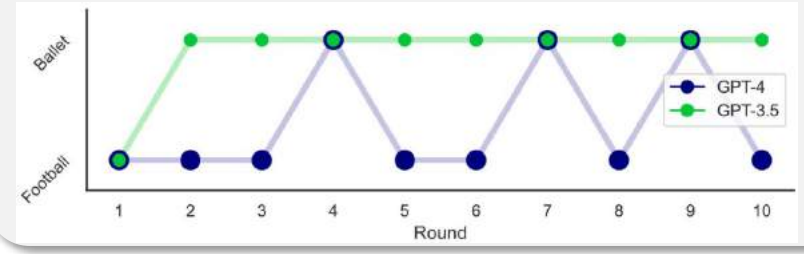

Reasoning

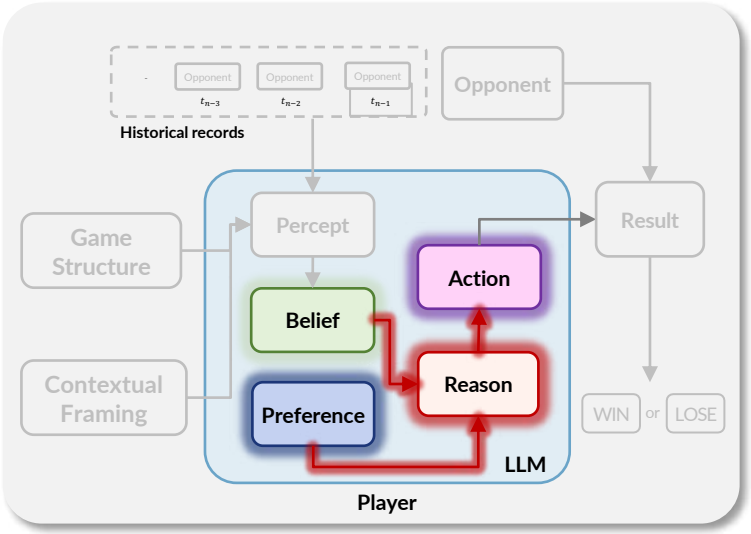

Can LLMs Reason based on Belief?

Implicit Belief

Player

Opponent	U	V
X	10, 0	0, 5
Y	5, 5	15, 10

Explicit Belief

Player

Opponent	U	V
X	0	5
Y	5	10

$a_o = Y$

Player

	U	V
X	10	0
Y	5	15

$D_o(Y,U) > D_o(X,U)$
$D_o(Y,V) > D_o(X,V)$

Given Belief

$D_m(V|Y) > D_m(U|Y)$

$a_m = V$

Player

	U	V
X	10	0
Y	5	15

(a)

Player

	U	V
X	8	7
Y	7	8

(b)

Player

	U	V
X	10	0
Y	5	6

(c)

Player

	U	V
X	40	0
Y	5	15

(d)

Belief

GPT-3.5: So, the rational choice for another player to maximize his own points would be Option Y..

$p(a_o|M)$

Action

GPT-3.5: Option U gives me the chance to win 40 points. … the most rational choice for me is to choose Option U.

$p(a_m|a_o, M)$

Belief

GPT-4: So, in summary, considering only their own point gain, the other player would choose Option Y.

$p(a_o|M)$

Action

GPT-4: … considering the possibility that the other player might switch to Option X, your best choice is Option U …

$p(a_m|a_o, M)$

Reasoning

Obvious

You are playing a game repeatedly with another player. In this game, you can choose between Option J and Option F. You will play 10 rounds in total with the same player. **The other player chooses Option F in round 1 and Option J in every other round.**

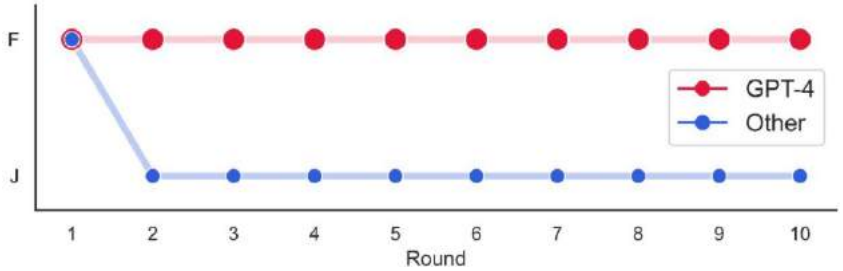

Beware of mistakes

You are playing a game repeatedly with another player. In this game, you can choose between Option J and Option F. You will play 10 rounds in total with the same player. **Be aware that the other player can make mistakes sometimes.**

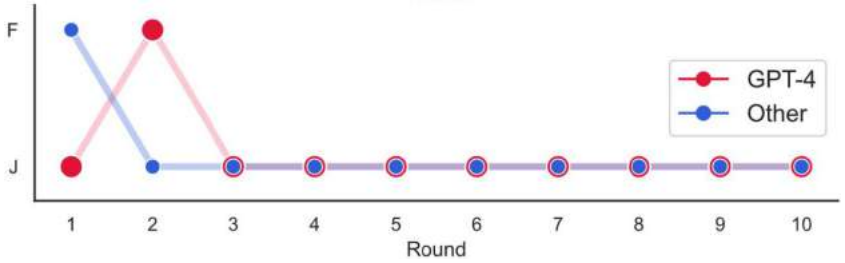

LLMs do not have the ability to autonomously follow human behavior in the game process. As a result, it is necessary to explicitly decouple human behavior for LLMs in game theory. However, even in the explicit game process, LLMs still appear to overlook / modify the refined belief. One possible solution is to transform the refined belief into the given belief in the dialogue.

Reasoning : Theory-of-Mind

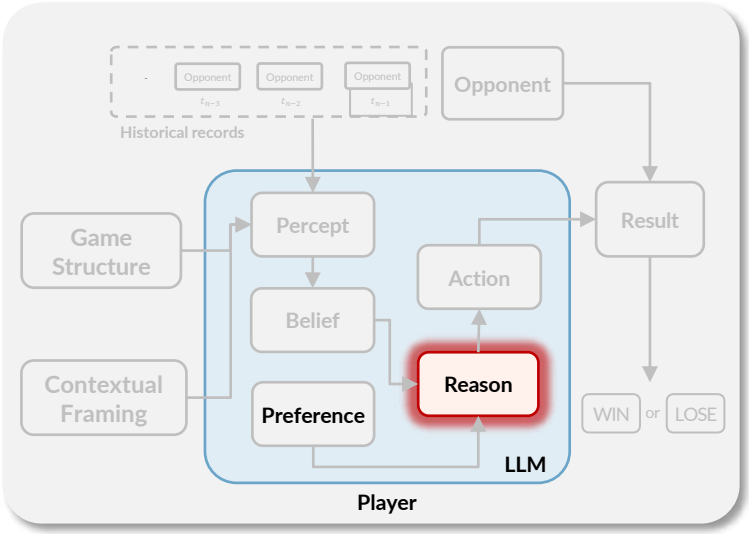

First-Order ToM Modelling

From my perspective, please infer several beliefs about the opponent's game pattern/preference for each round when holding different cards and the public card (if have).

Second-Order ToM Modelling

From my perspective, please infer under what circumstances is the opponent likely to be influenced by my actions? Additionally, in what situations would the opponent make decisions based solely on their own hand? From the perspective of the opponent (he cannot observe my card but only action), please infer several beliefs about my game pattern/preference when holding different cards.

The theory of mind (ToM) can enhance GPT's performance in imperfect information games.

Neural Theory-of-Mind

Jan 1, 2023

Minding Language Models' (Lack of) Theory of Mind: A Plug-and-Play Multi-Character Belief Tracker

Yejin Choi

Apr 3, 2023

Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs

Yejin Choi

"Static" Text

- reporting bias
- Lack of communicative intent and alternatives.
- Centering theory.

May 24, 2023

Clever Hans or Neural Theory of Mind? Stress Testing Social Reasoning in Large Language Models

Yejin Choi

Oct 22, 2023

Theory of Mind for Multi-Agent Collaboration via Large Language Models

University of Pittsburgh

Carnegie Mellon University, Pittsburgh

Evidence of **emergent** collaborative behaviors and high-order Theory of Mind capabilities among LLM-based agents.

Feb 23, 2024

8 tasks and 31 abilities in social cognition

ToMBENCH: Benchmarking Theory of Mind in Large Language Models

Tsinghua University

Feb 14, 2024

Longer and clearer narrative

Explicit personality **traits**

OpenToM: A Comprehensive Benchmark for Evaluating Theory-of-Mind Reasoning Capabilities of Large Language Models

King's College London

Huawei London Research Centre

The Alan Turing Institute

Oct 31, 2023

Interactions

Our results indicate that this capacity has **not yet emerged** in any manner.

FANTOM: A Benchmark for Stress-testing Machine Theory of Mind in Interactions

Yejin Choi

Oct 25, 2023

Higher-Order ToM

HI-TOM: A Benchmark for Evaluating Higher-Order Theory of Mind Reasoning in Large Language Models

University of Michigan

Westlake University

Advanced Methods

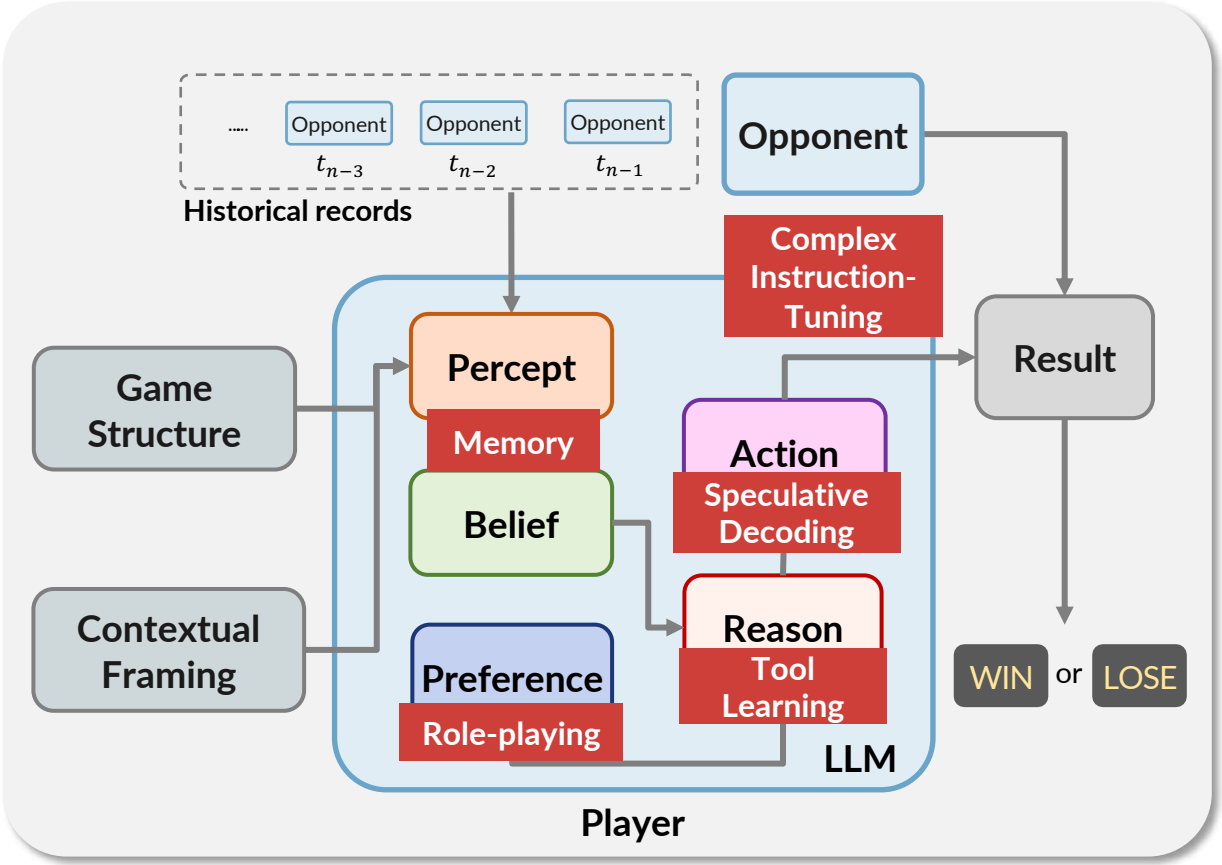

Modified based on *Can Large Language Models Serve as Rational Players in Game Theory? A Systematic Analysis*

Future Direction

Unified Framework

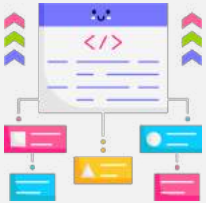

Unified Metrics

Massive Experiments

Complex Scenarios

Capability Enhancement

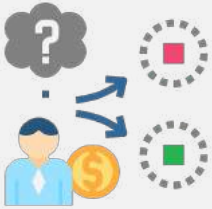

Practical Applications

Takeaway

- Behavioral science for machines is of vital importance.
- Existing research utilizes game theory as a theoretical framework to investigate the strategic reasoning capabilities of large language models (LLMs).
- Preliminary experimental results indicate that while current LLMs possess some strategic reasoning abilities, these capabilities are not consistently stable.
- AI researchers and social science researchers need to communicate more frequently to enhance the depth of their research, including AI for Social Science and Social Science of AI.

Thanks & QA

Xiachong Feng
Postdoc Fellow
HKU

Haochuan Wang
Intern
HKU

Lingpeng Kong
Assistant Professor
HKU

Chuan Wu
Professor
HKU